arXiv:1409.7074v1 [cond-mat.stat-mech] 24 Sep 2014

Variational Pseudolikelihood for Regularized Ising Inference

Charles K. Fisher
*Deptartment of Physics, Boston University, Boston, MA 02215**

I propose a variational approach to maximum pseudolikelihood inference of the Ising model. The variational algorithm is more computationally efficient, and does a better job predicting out-of-sample correlations than L_2 regularized maximum pseudolikelihood inference as well as mean field and isolated spin pair approximations with pseudocount regularization. The key to the approach is a variational energy that regularizes the inference problem by shrinking the couplings towards zero, while still allowing some large couplings to explain strong correlations. The utility of the variational pseudolikelihood approach is illustrated by training an Ising model to represent the letters A-J using samples of letters from different computer fonts.

Statistical mechanical models constructed from experimental observations provide a valuable tool for studying complex systems. The utility of the statistical mechanics approach to data-driven modeling is especially apparent in biology, providing insights into the behavior of flocking birds [1], the organization of neural networks in the brain [2–4], the structure and evolution of proteins [5–9], and many other topics [10, 11]. Generally, the 'inverse' statistical mechanics approach refers to the construction of statistical models using the principle of maximum entropy [12, 13]. In this approach, one constructs the probability distribution with the maximum entropy subject to constraints on its moments, which are derived from observations. Although maximum entropy modeling has been applied successfully to many different problems, it is still not clear how to estimate the parameters of the resulting probability distribution in an optimal way.

In this work, I consider the problem of estimating the parameters of an Ising model (see [14] for a review). The Ising model describes the statistics of a vector $\vec{\sigma}$ of N spin variables $\sigma_i \in \{-1, +1\}$, and can be derived from the principle of maximum entropy with constraints on the moments $\langle\sigma_i\rangle$ and $\langle\sigma_i\sigma_j\rangle$. The probability distribution for $\vec{\sigma}$ is given by $P(\vec{\sigma}) = Z^{-1}\exp(-U(\vec{\sigma}))$ where the energy is $U(\vec{\sigma}) = -\sum_i h_i\sigma_i - \sum_{i<j} J_{ij}\sigma_i\sigma_j$ and Z is a normalization constant. Here, h_i is a local field that biases σ_i towards the $+1$ or -1 configuration, and J_{ij} is a coupling that describes the strength of the interaction between spins i and j. The goal of the inverse Ising problem is to infer the parameters, i.e. h_i and J_{ij}, from a set of n observed configurations $\vec{\sigma}^{(l)}$ for $l = 1 \ldots n$.

The Ising model presents two obstacles that make the inverse problem quite difficult. First, computing the partition function (or its derivatives) is computationally intractable for large spin systems. As a result, the inverse Ising problem has to be solved approximately. A number of approximate methods for inferring the parameters of the Ising model have been introduced including naive mean field theory [15], the Thouless-Anderson-Palmer (TAP) approximation [16], the isolated spin pair approximation [14], the Sessak-Monasson expansion [17], and others [5, 18–27]. The second obstacle – overfitting – is more fundamental and generally affects all high dimensional problems in statistical inference.

Overfitting is a simple concept to illustrate for the Ising model. All of the moments, i.e. $\langle\sigma_i\rangle$ and $\langle\sigma_i\sigma_j\rangle$, are noisy because they are computed from a finite sample of n observed configurations. Thus, fitting all of the moments exactly necessarily incorporates noise from the finite sample size. As a consequence, parameters obtained by fitting the moments of one dataset may not provide a good description of a new dataset derived from the same system. In general, overfitting becomes a serious problem when the number of parameters (i.e. N^2) exceeds the number of independent observations (i.e. n).

A common approach to mitigating the effects of overfitting is to penalize parameters with large values [28]. For example, the Ising model can be 'regularized' by adding an L_2 penalty $\lambda\sum_{i<j} J_{ij}^2$ (as in [28]) or an L_1 penalty $\lambda\sum_{i<j}|J_{ij}|$ (as in [18, 21]) to the objective function that describes the fit to the data. Alternatively, the empirical moments can be modified using a 'pseudocount' according to the rules $\langle\sigma_i\rangle \to (1-\alpha)\langle\sigma_i\rangle$ and $\langle\sigma_i\sigma_j\rangle \to \delta_{ij} + (1-\delta_{ij})(1-\alpha)\langle\sigma_i\sigma_j\rangle$ [28]. The intuition for why regularization works comes from the law of total variance, which implies that the error in a statistical estimator is $error^2 = variance + bias^2$. Thus, a regularization method that decreases the variance in an estimator more than it increases the squared bias will have a smaller error and, therefore, better out-of-sample predictive ability.

The free parameter ($\lambda \geq 0$ or $0 \leq \alpha \leq 1$) that penalizes large couplings must be carefully chosen to ensure that the regularization method actually improves out-of-sample performance. Typically, the choice of the regularization parameter and subsequent testing of predictive ability are performed using cross-validation. To perform cross-validation, the observed data are randomly partitioned into three mutually exclusive datasets. The first dataset, usually called the 'training' sample, is used for fitting the model. Here, the training sample will be denoted 'F' for 'fitting'. The resulting parameters are used to predict the data contained in the second dataset, called the 'validation' sample ('V'). The optimal value

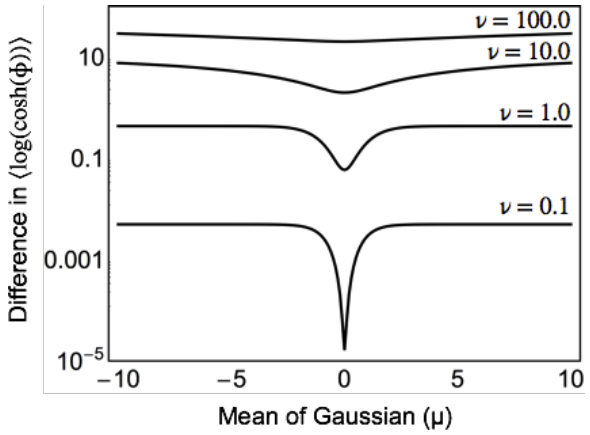

FIG. 1. A figure demonstrating that $\log(\cosh(\mu)) + \log(\cosh(\sigma)) - \langle\log(\cosh(\phi))\rangle$ (the vertical axis) is greater than zero when ϕ is Gaussian distributed with mean μ and standard deviation ν.

for the regularization parameter is chosen to maximize the predictive performance of the model on the validation sample. Finally, the out-of-sample performance of the resulting model is tested by measuring the agreement between the predictions of the model and the data contained in the third dataset, called the 'test' sample ('T').

Focusing on out-of-sample predictive ability provides a different metric than most previous papers on the inverse Ising model, which have studied the ability to reconstruct a coupling matrix from sampled configurations (e.g. [14]). In general, the performance of various algorithms for the inverse Ising model have been studied by defining a set of couplings, simulating configurations from the resulting Ising model, and then asking how well the couplings inferred by various methods reproduce the true J_{ij}'s. By contrast, this work will compare the performance of different approaches to the inverse Ising model by their out-of-sample predictive performance using a dataset of images developed for testing machine learning algorithms.

In order to compare the predictive performance of various inference methods for the Ising model, it is necessary to have a metric that quantifies the agreement between the model and the data in the test sample that can be easily computed. Here, I will use the negative log-pseudolikelihood, $\mathcal{L}_p^{data}(h,J) = -\langle\log\prod_i P(\sigma_i|\sigma_{j\neq i})\rangle_{data}$, where the angular brackets denote an average over the spin configurations in dataset 'F', 'V', or 'T'. Recently, Aurell and Ekeberg [18] demonstrated that the couplings of the Ising model can be inferred by minimizing the negative log-pseudolikelihood, which is given by (ignoring constant terms):

$$\mathcal{L}_p^F(h,J) = -\langle\log\prod_i P(\sigma_i|\sigma_{j\neq i})\rangle_F \tag{1}$$
$$= -\sum_i h_i\langle\sigma_i\rangle_F - \sum_{i,j\neq i} J_{ij}\langle\sigma_i\sigma_j\rangle_F + \sum_i\langle\log(\cosh(\phi_i))\rangle_F$$

where $\phi_i = h_i + \sum_{j\neq i} J_{ij}\sigma_j$ is the effective field acting on spin i due to the configurations of the other $n-1$ spins. The derivatives of $\mathcal{L}_p(h,J)$ are easy to calculate:

$$\frac{\partial\mathcal{L}_p^F}{\partial h_i} = \langle\sigma_i\rangle_F - \langle\tanh(\phi_i)\rangle_F$$
$$\frac{\partial\mathcal{L}_p^F}{\partial J_{ij}} = 2\langle\sigma_i\sigma_j\rangle_F - \langle v_i\tanh(\phi_j)\rangle_F - \langle v_j\tanh(\phi_i)\rangle_F$$

Note that these expressions assume $J_{ii} = 0$ and $J_{ij} = J_{ji}$. Thus, the Ising model can be fit to the training sample by minimizing $\mathcal{L}_p^F(h,J)$ using gradient descent.

While the computing the pseudolikelihood is certainly tractable it is, neverthless, computationally challenging to compute the averages over the training sample at every step of gradient descent if the sample size is large. Therefore, I propose a variational approach to maximum pseudolikelihood inference of the Ising model that is more computationally efficient and, as I will show, performs better out of sample. The variational pseudolikelihood approach minimizes an energy that provides an upper bound on $\mathcal{L}_p^F(h,J)$, given by:

$$\mathcal{E}(h,J) = -\sum_i h_i\langle\sigma_i\rangle_F - \sum_i\sum_{j\neq i} J_{ij}\langle\sigma_i\sigma_j\rangle_F$$
$$+\sum_i(\log(\cosh(\mu_i)) + \log(\cosh(\nu_i))) \tag{2}$$

The upper bound follows from $\langle\log(\cosh(\phi_i))\rangle \leq \log(\cosh(\mu_i)) + \log(\cosh(\nu_i))$ if ϕ_i is Gaussian distributed with mean μ_i and standard deviation ν_i; this inequality is illustrated in Fig. 1. Since ϕ_i is a sum of many random variables, I will assume that it is approximately Gaussian distributed with mean $\mu_i = h_i + \sum_{j\neq i} J_{ij}m_j$ and variance $\nu_i^2 = \sum_{j\neq i}\sum_{k\neq i} J_{ij}J_{ik}C_{jk}$, where $m_i = \langle s_i\rangle_F$, $C_{ij} = \langle s_is_j\rangle_F - \langle s_i\rangle_F\langle s_j\rangle_F$, and angular brackets denote averages over the training sample 'F'. Minimizing the variational energy over h_i yields the estimate

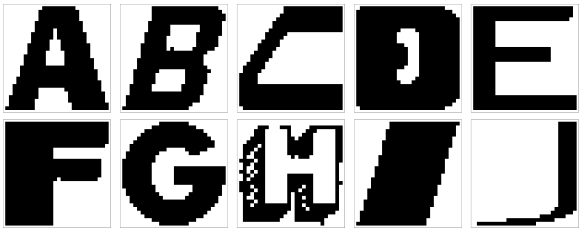

FIG. 2. Randomly chosen black and white images of letters A-J from the letters dataset.

TABLE I. Comparison of different methods for Ising inference according to their negative log-pseudolikelihoods on a dataset of letters. A more negative value indicates a better fit. The best performing methods for the training sample and the test sample are indicated by bold font. Pseudolikelihood and Variational Pseudolikelihood were fit using gradient descent with $\epsilon = 0.01$ and $\rho = 0.5$ for 10000 steps. In all cases, the local fields were estimated using $\hat{h}_i = \tanh^{-1}(m_i) - \sum_{j \neq i} J_{ij} m_j$.

Method	Training Sample	Test Sample
Variational Pseudolikelihood	- 475.2	**- 445.7**
Pseudolikelihood	**-535.8**	-303.3
Pseudolikelihood + L_2	-414.3	-398.9
Naive Mean Field (NMF)	-443.1	-282.5
NMF with pseudocount	- 423.8	-310.4
Isolated Spin Pair (ISP)	44.7	46.5
ISP with pseudocount	-109.2	-108.1

$\hat{h}_i = \tanh^{-1}(m_i) - \sum_{j \neq i} J_{ij} m_j$. After plugging in the estimate for h_i, it is only necessary to minimize a function of J given by:

$$G(J) = -\sum_i \sum_{j \neq i} J_{ij} C_{ij} + \sum_i \log(\cosh(\nu_i)) \qquad (3)$$

The derivatives are given by:

$$\frac{\partial G}{\partial J_{ij}} = (DJC)_{ij} + (DJC)_{ji} - 2C_{ij} = G'_{ij} \qquad (4)$$

where $D_{ij} = \tanh(\nu_i)\delta_{ij}/\nu_i$ and the diagonal elements of J are constrained to $J_{ii} = 0$. The couplings can be estimated by minimizing Eq. 3 using the gradient descent update $J_{ij}(t+1) = J_{ij}(t) - \epsilon(G'_{ij}(t) + \rho G'_{ij}(t-1))$, where ϵ is a small step size and ρ is a 'momentum'.

The variational pseudolikelihood approach was compared to ([14]): direct pseudolikelihood maximization (with and without an L_2 penalty on the couplings), naive mean field (NMF) inversion (with and without pseudocount regularization), and the isolated spin pair (ISP) approximation (with and without pseudocount regularization). For the regularized methods, the penalty parameters were chosen by cross-validation to minimize $\mathcal{L}_p^V(h, J)$ on the validation sample. The performances of these methods was assessed using training, validation, and test samples con structured from the 'notMNIST' dataset, which consists of 28×28 pixel images of the letters A-J compiled from a large number of different fonts (see Appendix). The 'notMNIST' dataset contains roughly 50000 images for each letter, and the samples of each letter are very diverse. Examples of binarized images for letters A-J are shown in Fig. 2. Mutually exclusive training, validation, and test samples were constructed, each containing 500 randomly chosen images for each letter.

Table I shows a comparision of these methods according to $\mathcal{L}_p^F(h, J)$ for the training set images, and $\mathcal{L}_p^T(h, J)$ for the test set images. As one would expect, ranking by the pseudolikelihood ensures that direct pseudolikelihood maximization is the best performing method on the training data. However, the variational pseudolikelihood approach presented here outperforms all other methods, including direct pseudolikelihood maximization, when compared by out-of-sample performance on the test set. In addition, a simple implementation of variational pseudolikelihood inference was approximately an order of magnitude faster than direct pseudolikelihood maximization running on a desktop computer, even ignoring the time spent choosing the penalty parameter for L_2 regularized pseudolikelihood.

Comparing the out-of-sample results for the unregularized and regularized versions of maximum pseudolikelihood, NMF, and ISP estimates in Table I demonstrates the importance of regularization for the inverse Ising model. Histograms of the couplings inferred with direct pseudolikelihood maximization, pseudlikelihood maximization with an L_2 penalty, and the variational pseudolikelihood method (Fig. 2) illustrate that the methods which performed better out-of-sample had smaller estimates for the couplings. However, adding the L_2 penalty to the pseudolikelihood resulted in an inferred Ising model with many weak couplings, and no strong couplings. The variational approach, by contrast, produced both strong and weak couplings, which resulted in improved out-of-sample prediction performance relative to other methods.

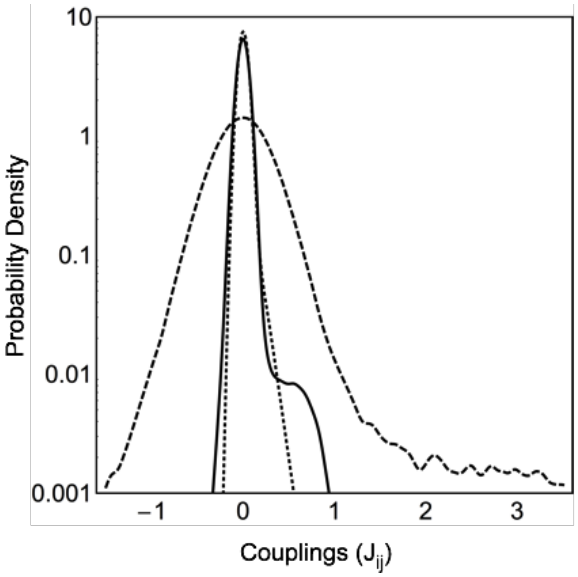

FIG. 3. Smooth histograms of the inferred couplings obtained from direct pseudolikelihood maximization (dashed line), pseudolikelihood maximization with an L_2 penalty (dotted line), and the variational pseudolikelihood method (solid line).

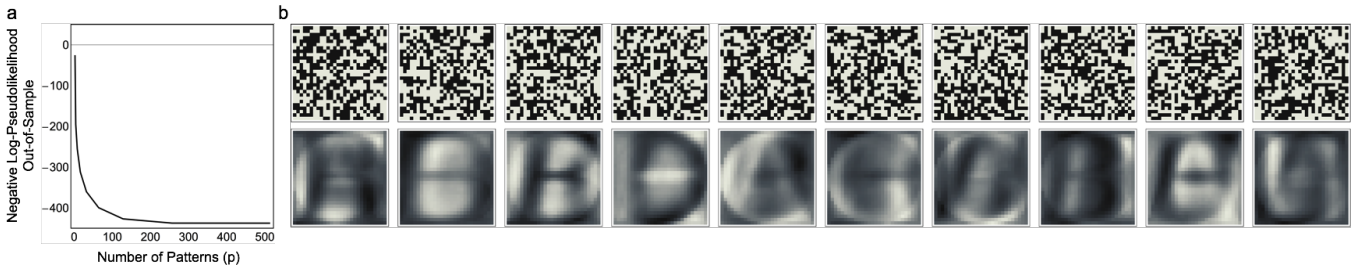

FIG. 4. Inference of Hopfield-like patterns using variational pseudolikelihood. a) Negative log-pseudolikelihood on the test sample as a function of the number of patterns. b) The top row shows the random initial conditions for the 10 patterns. The bottom row shows the 10 patterns obtained by gradient descent minimization of the variational pseudolikelihood with $\epsilon = 0.001$ and $\rho = 0.5$. The model achieved negative log-pseudolikelihoods of $\mathcal{L}_p^F(h, J) = -277.1$ and $\mathcal{L}_p^T(h, J) = -275.0$ on the training sample and the test sample, respectively.

The analysis presented above demonstrates that variational pseudolikelihood inference can be used to train an Ising model so that it makes good out-of-sample predictions. Here, I provide an example to illustrate how the variational pseudolikelihood method can be used in practice. I examined a simple question: what features does the Ising model learn when trained on images of the letters A-J? Using the variational pseudolikelihood approach, it is simple to infer coupling matrices with a specific structure, such as that of a Hopfield neural network [5, 21, 26, 29]. The couplings in a Hopfield network are described by a Hebbian rule such that $J_{ij} = \sum_s \xi_i^{(s)} \xi_j^{(s)} (1 - \delta_{ij})$, where $\vec{\xi}^{(s)}$ for $s = 1 \dots p$ are 'patterns' stored in the memory. These patterns describe features in the correlation structure of the data. The derivatives of the variational energy with respect to the patterns can be computed using the chain rule:

$$\frac{\partial G}{\partial \xi_i^{(s)}} \propto \sum_j G'_{ij} \xi_j^{(s)} \tag{5}$$

Thus, the patterns can be inferred quite easily by minimizing the variational energy using gradient descent. Note, however, that the patterns are unconstrained with respect to orthogonal transformations [30], and the symmetry is broken by the initial conditions of the optimization.

The letters in the notMNIST dataset are quiet diverse, and hundreds of patterns are required to approach the performance of the unconstrained variational pseudolikelihood method (Fig. 4a). Nevertheless, the $p = 10$ Hopfield model provides a useful visual example. Starting from randomly initialized patterns (shown in the top row of Fig. 4b), the patterns converge after less than 1000 steps of gradient descent to images with recognizable features of letters (shown in the second row of Fig. 4b).

The notMNIST images present a challenging dataset for the Ising model. State-of-the-art machine learning methods for modeling similar types of images include a variety of pre-processing and fine tuning steps [31], and the best performing algorithms are based on deep neural networks that are believed to excel at capturing higher order correlation structure in the data that cannot be described by the Ising model [32]. Nevertheless, Fig. 4 clearly demonstrates an Ising model based on a few Hopfield patterns is able to capture recognizable features from the data.

In summary, I presented a variational approach to maximum pseudolikelihood inference of the Ising model. The variational pseudolikelihood method outperformed all other tested methods at describing the out-of-sample correlation structure in a dataset consisting of images of the letters A-J from various computer fonts. This method has many characteristics that make it attractive in practice: it is fast, accurate, can be used with large spin systems, and can easily be adapted to learn coupling matrices with specific structures such as Hopfield models. Moreover, it should be easy to extend the approach presented here to construct variational pseudolikelihood algorithms for inferring Potts models [6], Restricted Boltzmann Machines [33], and other inverse models for spin systems.

I would like to thank Alex H. Lang, Thomas C. Lang, Pankaj Mehta, Ilya Nemenman, and Javad Noorbakhsh for helpful discussions. Funding for this research was provided by a Simons Investigator Award to Pankaj Mehta.

Appendix: The 'notMNIST' dataset consists of 28×28 pixel images of the letters A-J compiled from a large number of different fonts. The images are available at http://yaroslavvb.blogspot.com/2011/09/notmnist-dataset.html. Each image in the original dataset is grayscale, with pixels taking on values between 0 and 255. The images were binarized in the simplest way by defining pixels with intensity less than or equal to 127 as spin -1, and those with intensity greater than 127 as spin $+1$. Finally, each image was flattened into a vector of $N = 784$ spins.

* charleskennethfisher@gmail.com

[1] W. Bialek, A. Cavagna, I. Giardina, T. Mora, E. Silvestri, M. Viale, and A. M. Walczak, Proceedings of the National Academy of Sciences **109**, 4786 (2012).

[2] G. Tkačik, O. Marre, T. Mora, D. Amodei, M. J. Berry II, and W. Bialek, Journal of Statistical Mechanics: Theory and Experiment **2013**, P03011 (2013).

[3] S. Cocco and R. Monasson, BMC Neuroscience **12**, P224 (2011).

[4] Y. Roudi, J. Tyrcha, and J. Hertz, Physical Review E **79**, 051915 (2009).

[5] S. Cocco, R. Monasson, and M. Weigt, PLoS computational biology **9**, e1003176 (2013).

[6] M. Ekeberg, C. Lövkvist, Y. Lan, M. Weigt, and E. Aurell, Physical Review E **87**, 012707 (2013).

[7] M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, and T. Hwa, Proceedings of the National Academy of Sciences **106**, 67 (2009).

[8] T. A. Hopf, L. J. Colwell, R. Sheridan, B. Rost, C. Sander, and D. S. Marks, Cell **149**, 1607 (2012).

[9] A. Schug, M. Weigt, J. N. Onuchic, T. Hwa, and H. Szurmant, Proceedings of the National Academy of Sciences **106**, 22124 (2009).

[10] T. Mora, A. M. Walczak, W. Bialek, and C. G. Callan, Proceedings of the National Academy of Sciences **107**, 5405 (2010).

[11] M. Santolini, T. Mora, and V. Hakim, arXiv preprint arXiv:1302.4424 (2013).

[12] E. T. Jaynes, Proceedings of the IEEE **70**, 939 (1982).

[13] S. Pressé, K. Ghosh, J. Lee, and K. A. Dill, Reviews of Modern Physics **85**, 1115 (2013).

[14] Y. Roudi, E. Aurell, and J. A. Hertz, Frontiers in computational neuroscience **3** (2009).

[15] H. J. Kappen and F. Rodriguez, Neural Computation **10**, 1137 (1998).

[16] T. Tanaka, Physical Review E **58**, 2302 (1998).

[17] V. Sessak and R. Monasson, Journal of Physics A: Mathematical and Theoretical **42**, 055001 (2009).

[18] E. Aurell and M. Ekeberg, Physical review letters **108**, 090201 (2012).

[19] H. C. Nguyen and J. Berg, Journal of Statistical Mechanics: Theory and Experiment **2012**, P03004 (2012).

[20] I. Mastromatteo, Journal of Statistical Physics **150**, 658 (2013).

[21] H. Huang, The European Physical Journal B **86**, 1 (2013).

[22] S. Cocco and R. Monasson, Journal of Statistical Physics **147**, 252 (2012).

[23] H. C. Nguyen and J. Berg, Physical review letters **109**, 050602 (2012).

[24] H. Huang and Y. Kabashima, Physical Review E **87**, 062129 (2013).

[25] S. Cocco and R. Monasson, Physical review letters **106**, 090601 (2011).

[26] H. Huang, Physical Review E **81**, 036104 (2010).

[27] E. Aurell, C. Ollion, and Y. Roudi, The European Physical Journal B-Condensed Matter and Complex Systems **77**, 587 (2010).

[28] J. Barton, S. Cocco, E. De Leonardis, and R. Monasson, arXiv preprint arXiv:1405.0233 (2014).

[29] J. J. Hopfield, Proceedings of the national academy of sciences **79**, 2554 (1982).

[30] S. Cocco, R. Monasson, and M. Weigt, in *Journal of Physics: Conference Series*, Vol. 473 (IOP Publishing, 2013) p. 012010.

[31] Y. Tang and I. Sutskever, *Data normalization in the learning of restricted Boltzmann machines*, Tech. Rep. (Technical Report UTML-TR-11-2, Department of Computer Science, University of Toronto, 2011).

[32] D. Ciresan, U. Meier, and J. Schmidhuber, in *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on* (IEEE, 2012) pp. 3642–3649.

[33] R. Salakhutdinov, A. Mnih, and G. Hinton, in *Proceedings of the 24th international conference on Machine learning* (ACM, 2007) pp. 791–798.